arXiv:1603.06185v3 [cs.IT] 24 Jul 2016

Optimizing Data Freshness, Throughput, and Delay in Multi-Server Information-Update Systems

Ahmed M. Bedewy†, Yin Sun†, and Ness B. Shroff† ‡

†Dept. of ECE, ‡Dept. of CSE, The Ohio State University, Columbus, OH.

***Abstract*—In this work, we investigate the design of information-update systems, where incoming update packets are forwarded to a remote destination through multiple servers (each server can be viewed as a wireless channel). One important performance metric of these systems is *the data freshness at the destination*, also called the *age-of-information* or simply *age*, which is defined as the time elapsed since the freshest packet at the destination was generated. Recent studies on information-update systems have shown that the age-of-information can be reduced by intelligently dropping stale packets. However, packet dropping may not be appropriate in many applications, such as news and social updates, where users are interested in not just the latest updates, but also past news. Therefore, all packets may need to be successfully delivered. In this paper, we study how to optimize age-of-information without throughput loss. We consider a general scenario where incoming update packets do not necessarily arrive in the order of their generation times. We prove that a preemptive Last Generated First Served (LGFS) policy simultaneous optimizes the age, throughput, and delay performance in infinite buffer queueing systems. We also show age-optimality for the LGFS policy for any finite queue size. These results hold for arbitrary, including non-stationary, arrival processes.**

I. Introduction

The ubiquity of mobile devices and applications, has increased the demand for real-time information updates, such as news, weather reports, email notifications, stock quotes, social updates, mobile ads, etc. Also, in network-based monitoring and control systems, timely status updates are crucial. These include, but are not limited to, sensor networks used in temperature or other physical phenomenon, and autonomous vehicle systems.

A common objective in these applications is to keep the destination updated with the latest information. To identify the timeliness of the updates, a metric called *data freshness*, also called *age of information*, or simply *age*, was defined in [1]–[4]. At time t, if $U(t)$ is the time when the freshest update at the destination was generated, age $\Delta(t)$ is $\Delta(t) = t - U(t)$. Hence, age is the time elapsed since the freshest packet was generated.

There have been several recent works on characterizing the time-average age of different information-update policies under Poisson arrival process, and finding policies with a small time-average age [4]–[10]. In [4]–[6], the update generation rate was optimized to improve data freshness in First-Come First-Served (FCFS) information-update systems. To improve the age, these studies also reduced the update generation rate, which in turn sacrificed the system throughput. In [7], [8], it was found that the age can be improved by discarding old packets waiting in the queue if a new sample arrives. This can greatly reduce the impact of queueing delay on data freshness. However, many applications may not want to discard packets, e.g., where the users are interested in not just the latest updates, but also past news, in which case all packets must be successfully delivered. In [9], [10], the time-average age was characterized for Last-Come First-Served (LCFS) information-update systems with and without preemption; and FCFS with two servers under a Poisson arrival process. Applications of information updates in channel information feedback and sensor networks were considered in [11]–[13].

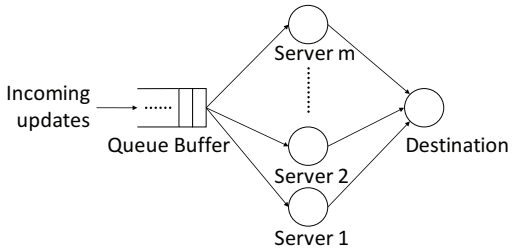

Figure 1: System model.

Another important problem is how to maximize data freshness in information-update systems. This involves jointly controlling both the generation and transmission of packet updates [12]–[14]. An information update policy was developed in [14], which was proven to minimize the time-average age and time-average age penalty among all causally feasible policies. In this setting, a counter-intuitive phenomenon was revealed: While a zero-wait or work-conserving policy, that generates and submits a fresh update once the server becomes idle, achieves the maximum throughput and the minimum average delay, surprisingly, this zero-wait policy does not always minimize the age. This implies that there is no policy that can simultaneously minimize age and maximize throughput, if the generation and transmission of update packets are jointly controlled.

In this paper, we consider an information-update system which enqueues incoming update packets and forwards them to a remote destination through multiple servers, as shown in Fig. 1. In this setting, the updates are generated exogenously to the system, which is different from [14]. We aim to answer the following questions: How to establish age-optimality in a general policy space and under arbitrary arrival process? Is

it possible to simultaneously optimize multiple performance metrics, such as age, throughput, and delay? To that end, the following are the key contributions of this paper:

- We consider a general scenario where the update packets do not necessarily arrive in the order of their generation times, which has not been considered before. We prove that, if the packet service times are *i.i.d.* exponentially distributed, then for an arbitrary arrival process and any queue size, a preemptive Last-Generated First-Served (LGFS) policy achieves an age process that is stochastically smaller than any causally feasible policies (Theorem 1). This implies that the preemptive LGFS policy minimizes many data freshness metrics, including time-average age, average peak age, and time-average age penalty (Corollary 6). The intuition is that the freshest update packets are served as early as possible in the preemptive LGFS policy. In particular, the distribution of the age process of the preemptive LGFS policy is invariant over all queue sizes.
- In addition, we show that if the buffer has an infinite size, then the preemptive LGFS policy is also throughput-optimal and delay-optimal among all causally feasible policies (Theorem 7). To the best of our knowledge, this is the first study which simultaneously optimizes data freshness, throughput, and delay in information-update systems.

We note that when the incoming update packets are arriving in the same order of their generation times, the proposed LGFS policy is identical to the LCFS policy studied in [9]. In particular, the time-average age of preemptive and non-preemptive LCFS policies are analyzed in [9] for single-server queueing systems with Poisson arrival process and a queue size of one packet. This paper complements and generalizes the results in [9] by (i) allowing the incoming updates to not arrive in the order of their generation times, (ii) considering more general multi-server queueing systems with arbitrary update arrivals and arbitrary queue size, and (iii) providing an age-optimality proof.

This paper is motivated by the applications that require all the updates to be sent not only the latest one. For instance, it was shown that the completely stale channel state information (CSI) is still very useful in a MIMO broadcast channel in [15]. Particularly, by using not only the latest CSI but also the outdated CSI, one can achieve non-trivial multiplexing gains. Another example is stock predictions. In [16], a base stock policy was introduced such that it uses stale or old forecasts beside the current forecast to set base stock levels at the manufacturer. It was found that, in case of a non-stationary demand process, the proposed policy can decrease the expected supply chain inventory and shortage costs; and significantly reduce the fluctuations in production levels compared to that of using current information.

Another important application for our model is the reinforcement learning. The reinforcement learning is learning what action should be performed in order to maximize a numerical reward signal [17]. The reinforcement learning agent is not told which action to take, but instead it must figure out which actions that leads to maximum reward by trying them. If the learning agent takes action depending on the updates received from a certain source, then its decision requires the knowledge of all updates in addition to the latest one. Another application is dynamic-content distribution service. For example, a group of wireless users receives frequent content updates from a common source such as news feeds. In this case, the user is interested in not only the latest news, but also the previous news-feeds.

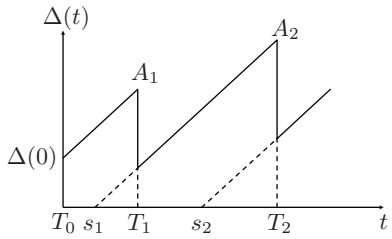

Figure 2: Evolution of the age-of-information $\Delta(t)$.

II. System Model

We consider an information-update system depicted in Fig. 1, where an infinite sequence of update packets are sent to a destination through m identical servers. Each server could be a wireless channel, a TCP connection, etc. The update packets are generated at time $s_1, s_2, \ldots$, where $s_1 = 0 \leq s_2 \leq \ldots$ are arbitrarily given. These packets do not need to arrive to the system in the order that they were generated. Let a_i be the arrival time of the packet generated at time s_i, such that $s_i \leq a_i$. As the packets may arrive out of order, it may happen that $a_i > a_{i+1}$. Hence, each packet is identified by the pair (s_i, a_i), which will be used to make scheduling decision. After arrival, each update packet is stored in a queue, waiting to be assigned to one of the servers. Let B denote the buffer size of the queue which can be infinite, finite, or even zero. If B is finite, the queue buffer may overflow and some packets are dropped, which would incur a throughput loss. The packet service times are exponentially distributed with rate μ, which are *i.i.d.* across time and servers. Let T_i denote the i-th packet delivery time instant.

The system starts to operate at time $t = 0$. Define $U(t)$ as the time when the freshest update at the destination at time t was generated, where we set $U(t) = 0$ at $t = 0$. The *age-of-information*, or simply the *age*, is defined as

$$\Delta(t) = t - U(t). \tag{1}$$

From this definition, we can deduce that the age increases linearly with t but is reset to a smaller value with the delivery of each update that contains fresher information, as shown in Fig. 2. Define A_k as the k-th peak value of $\Delta(t)$ since time $t = 0$.

Following [14], we define an age penalty function $g(\Delta(t))$ to represent the level of "dissatisfaction" for data staleness or

the "need" for new information update, where $g(\cdot)$ is a general non-decreasing function, that is determined based on specific applications.

A. Scheduling Policy

A scheduling policy, denoted by π, assigns update packets to the servers over time. We first define several classes of policies:

Definition 1. Service Preemption: In a *preemptive* policy, a server can switch to send any packet at any time; the preempted packets will be stored back into the queue if there is enough buffer space and sent at a later time when the servers are available again. In contrast, in a *non-preemptive* policy, a server must complete delivering the current packet before starting to send another packet.

Definition 2. Work Conservation: A policy is said to be *work-conserving*, if no server is idle when there are packets waiting in the queue.

We use Π to denote the set of all causal policies, in which scheduling decisions are made based on the history and current state of the system, as well as the distribution of the packet service time. Let $\Pi_{wc} \subseteq \Pi$ be the set of work-conserving feasible policies, and $\Pi_{nwc} \subseteq \Pi$ be the set of non-work-conserving feasible policies.

B. Age Optimality

We will need the following definitions: Let $\vec{x} = (x_1, x_2, \ldots, x_n)$ and $\vec{y} = (y_1, y_2, \ldots, y_n)$ be two vectors in $\mathbb{R}^n$, then we denote $\vec{x} \leq \vec{y}$ if $x_i \leq y_i$ for $i = 1, 2, \ldots, n$.

Definition 3. Stochastic Ordering: [18] Let X and Y be two random variables. Then, X is said to be stochastically smaller than Y (denoted as $X \leq_{\text{st}} Y$), if

$$P\{X > x\} \leq P\{Y > x\}, \quad \forall x \in \mathbb{R}.$$

Definition 4. Multivariate Stochastic Ordering: [18] A set $U \subseteq \mathbb{R}^n$ is called upper if $\vec{y} \in U$ whenever $\vec{y} \geq \vec{x}$ and $\vec{x} \in U$. Let $\vec{X}$ and $\vec{Y}$ be two random vectors. Then, $\vec{X}$ is said to be stochastically smaller than $\vec{Y}$ (denoted as $\vec{X} \leq_{\text{st}} \vec{Y}$), if

$$P\{\vec{X} \in U\} \leq P\{\vec{Y} \in U\}, \quad \text{for all upper sets} \quad U \subseteq \mathbb{R}^n.$$

Definition 5. Stochastic Ordering of Stochastic Processes: [18] Let $\{X(t), t \in [0,\infty)\}$ and $\{Y(t), t \in [0,\infty)\}$ be two stochastic processes. Then, $\{X(t), t \in [0,\infty)\}$ is said to be stochastically smaller than $\{Y(t), t \in [0,\infty)\}$ (denoted by $\{X(t), t \in [0,\infty)\} \leq_{\text{st}} \{Y(t), t \in [0,\infty)\}$), if, for all choices of an integer n and $t_1 < t_2 < \ldots < t_n$ in $[0,\infty)$, it holds that

$$(X(t_1), X(t_2), \ldots, X(t_n)) \leq_{\text{st}} (Y(t_1), Y(t_2), \ldots, Y(t_n)), \quad (2)$$

where the multivariate stochastic ordering in (2) was defined in Definition 4.

Definition 6. Age Optimality: A policy $P \in \Pi$ is said to be *age-optimal*, if for all $\pi \in \Pi$

$$\{\Delta_p(t), t \in [0,\infty)\} \leq_{\text{st}} \{\Delta_\pi(t), t \in [0,\infty)\}. \quad (3)$$

As we will see in Corollary 6, the results obtained using this definition of age optimality is quite strong.

III. Optimality Analysis

Algorithm 1: Preemptive Last Generated First Served policy.

```
α := 0;
while the system is ON do
    if a new packet with time stamp s arrives then
        if s ≤ α then                    // The packet is outdated.
            Store the packet in the queue;
        else                             // The packet carries fresh information.
            if all servers are busy then
                The new packet is assigned to a server by
                preempting the packet with time stamp α;
                The preempted packet with time stamp α
                is stored back to the queue;
                Set α as the smallest time stamp of the
                packets under service;
            else                         // At least one of the servers is idle.
                Assign the new packet to one idle server;
            end
        end
    end
    if a packet is delivered then
        if the queue is not empty then
            Pick any packet in the queue and assign it to
            the idle server;
        end
    end
end
```

In this section, we study a LGFS policy, in which the packets under service are the freshest (i.e., are generated the latest) among all packets in the queue. The implementation details of a preemptive LGFS policy is depicted in Algorithm 1, where α is the smallest time stamp of the packets under service. We next show that the preemptive LGFS policy is age-optimal.

Theorem 1. If the packet service times are exponentially distributed and *i.i.d.* across time and servers, then for any given packet generation and arrival times, and any given queue size B, the preemptive LGFS policy is age-optimal.

We need to define the system state of any policy π:

Definition 7. At any time t, the system state of policy π is specified by $V_\pi(t) = (U_\pi(t), \alpha_{1,\pi}(t), \ldots, \alpha_{m,\pi}(t))$, where $U_\pi(t)$ is the largest time stamp of the packets that have already been delivered to the destination. Define $\alpha_{i,\pi}(t)$ as the i-th smallest time stamp of the packets being processed by the

servers. Without loss of generality, if k servers are sending stale packets (i.e., $\alpha_{1,\pi}(t) \leq \ldots \leq \alpha_{k,\pi}(t) \leq U_\pi(t)$) or k servers are idle, then we set $\alpha_{1,\pi}(t) = \ldots = \alpha_{k,\pi}(t) = U_\pi(t)$. Hence,

$$U_\pi(t) \leq \alpha_{1,\pi}(t) \leq \ldots \leq \alpha_{m,\pi}(t). \tag{4}$$

Let $\{V_\pi(t), t \in [0,\infty)\}$ be the state process of policy π, which is assumed to be right-continuous.

For notational simplicity, let policy P represent the preemptive LGFS policy. Then, by the construction of policy P, $\alpha_{1,P}(t), \alpha_{2,P}(t), \ldots, \alpha_{m,P}(t)$ are the time stamps of m freshest packets among all packets arrived during $[0,t]$.

The key step in the proof of Theorem 1 is the following lemma, where we compare policy P with any work-conserving policy $\pi \in \Pi_{wc}$.

Lemma 2. For any given packet generation and arrival times, and any given queue size B, suppose that $V_P(0^-) = V_\pi(0^-)$ for all work conserving policies $\pi \in \Pi_{wc}$, then

$$\{V_P(t), t \in [0,\infty)\} \geq_{\text{st}} \{V_\pi(t), t \in [0,\infty)\}. \tag{5}$$

We use coupling and forward induction to prove Lemma 2. For any work-conserving policy π, suppose that stochastic processes $\hat{V}_P(t)$ and $\hat{V}_\pi(t)$ have the same stochastic laws as $V_P(t)$ and $V_\pi(t)$. The packet deliveries of the state processes $\hat{V}_P(t)$ and $\hat{V}_\pi(t)$ are coupled in the following manner: If the packet with time stamp $\hat{\alpha}_{i,P}(t)$ is delivered at time t as $\hat{V}_P(t)$ evolves, then the packet with time stamp $\hat{\alpha}_{i,\pi}(t)$ is delivered at time t as $\hat{V}_\pi(t)$ evolves. Such a coupling is valid since the service time is exponentially distributed and thus memoryless. On the other hand, policy P and policy π have identical arrival processes. According to Theorem 6.B.30 in [18], if we can show

$$P\left[\hat{V}_P(t) \geq \hat{V}_\pi(t), t \in [0,\infty)\right] = 1, \tag{6}$$

then (5) is proven. Next, we use the following lemmas to prove (6):

Lemma 2*. Suppose that the system state of policy P is $\{\hat{U}_P, \hat{\alpha}_{1,P}, \ldots, \hat{\alpha}_{m,P}\}$, and meanwhile the system state of policy π is $\{\hat{U}_\pi, \hat{\alpha}_{1,\pi}, \ldots, \hat{\alpha}_{m,\pi}\}$. If

$$\hat{U}_P \geq \hat{U}_\pi, \tag{7}$$

then,

$$\hat{\alpha}_{i,P} \geq \hat{\alpha}_{i,\pi}, \quad \forall i = 1, \ldots, m. \tag{8}$$

Proof: See Appendix A. ∎

Lemma 3. Suppose that under policy P, $\{\hat{U}'_P, \hat{\alpha}'_{1,P}, \ldots, \hat{\alpha}'_{m,P}\}$ is obtained by delivering a packet with time stamp $\hat{\alpha}_{l,P}$ to the destination in the system whose state is $\{\hat{U}_P, \hat{\alpha}_{1,P}, \ldots, \hat{\alpha}_{m,P}\}$. Further, suppose that under policy π, $\{\hat{U}'_\pi, \hat{\alpha}'_{1,\pi}, \ldots, \hat{\alpha}'_{m,\pi}\}$ is obtained by delivering a packet with time stamp $\hat{\alpha}_{l,\pi}$ to the destination in the system whose state is $\{\hat{U}_\pi, \hat{\alpha}_{1,\pi}, \ldots, \hat{\alpha}_{m,\pi}\}$. If

$$\hat{\alpha}_{i,P} \geq \hat{\alpha}_{i,\pi}, \quad \forall i = 1, \ldots, m, \tag{9}$$

then,

$$\hat{U}'_P \geq \hat{U}'_\pi, \hat{\alpha}'_{i,P} \geq \hat{\alpha}'_{i,\pi}, \quad \forall i = 1, \ldots, m. \tag{10}$$

Proof: See Appendix B. ∎

Lemma 4. Suppose that under policy P, $\{\hat{U}'_P, \hat{\alpha}'_{1,P}, \ldots, \hat{\alpha}'_{m,P}\}$ is obtained by adding a packet with time stamp s to the system whose state is $\{\hat{U}_P, \hat{\alpha}_{1,P}, \ldots, \hat{\alpha}_{m,P}\}$. Further, suppose that under policy π, $\{\hat{U}'_\pi, \hat{\alpha}'_{1,\pi}, \ldots, \hat{\alpha}'_{m,\pi}\}$ is obtained by adding a packet with time stamp s to the system whose state is $\{\hat{U}_\pi, \hat{\alpha}_{1,\pi}, \ldots, \hat{\alpha}_{m,\pi}\}$. If

$$\hat{U}_P \geq \hat{U}_\pi, \tag{11}$$

then

$$\hat{U}'_P \geq \hat{U}'_\pi, \hat{\alpha}'_{i,P} \geq \hat{\alpha}'_{i,\pi}, \quad \forall i = 1, \ldots, m. \tag{12}$$

Proof: See Appendix C. ∎

Proof of Lemma 2: For any sample path, we have that $\hat{U}_P(0^-) = \hat{U}_\pi(0^-)$ and $\hat{\alpha}_{i,P}(0^-) = \hat{\alpha}_{i,\pi}(0^-)$ for $i = 1, \ldots, m$. This, together with Lemma 3 and 4, implies that

$$\hat{U}_P(t) \geq \hat{U}_\pi(t), \hat{\alpha}_{i,P}(t) \geq \hat{\alpha}_{i,\pi}(t),$$

holds for all $t \in [0,\infty)$ and $i = 1, \ldots, m$. Hence, (6) follows.

Because $\{\hat{V}_P(t), t \in [0,\infty)\}$ and $\{V_P(t), t \in [0,\infty)\}$ are of the same distribution, $\{\hat{V}_\pi(t), t \in [0,\infty)\}$ and $\{V_\pi(t), t \in [0,\infty)\}$ are of the same distribution, by Theorem 6.B.30 in [18] we get (5). This completes the proof. ∎

Proof of Theorem 1: As a result of Lemma 2, we have

$$\{U_P(t), t \in [0,\infty)\} \geq_{\text{st}} \{U_\pi(t), t \in [0,\infty)\}, \ \forall \pi \in \Pi_{wc},$$

which implies

$$\{\Delta_P(t), t \in [0,\infty)\} \leq_{\text{st}} \{\Delta_\pi(t), t \in [0,\infty)\}, \ \forall \pi \in \Pi_{wc}.$$

Finally, since the service times are *i.i.d.* across time and servers, service idling only increases the waiting time of the packet in the system. Therefore, the age under non-work-conserving policies will be greater. As a result, we have

$$\{\Delta_P(t), t \in [0,\infty)\} \leq_{\text{st}} \{\Delta_\pi(t), t \in [0,\infty)\}, \forall \pi \in \Pi.$$

This completes the proof. ∎

As a result of Theorem 1, we can deduce the following corollaries:

Corollary 5. If the packet service times are exponentially distributed and *i.i.d.* across time and servers, then for any given packet generation and arrival times, the age performance of the preemptive LGFS policy remains the same for any queue size $B \geq 0$.

Proof: In the proofs of Lemma 3 and Lemma 4, the system states evolution is invariant for any queue size $B \geq 0$. Hence, the age performance of the preemptive LGFS policy is invariant for any queue size $B \geq 0$. ∎

Corollary 6. If the packet service times are *i.i.d.* exponentially distributed across time and servers, then for any given packet

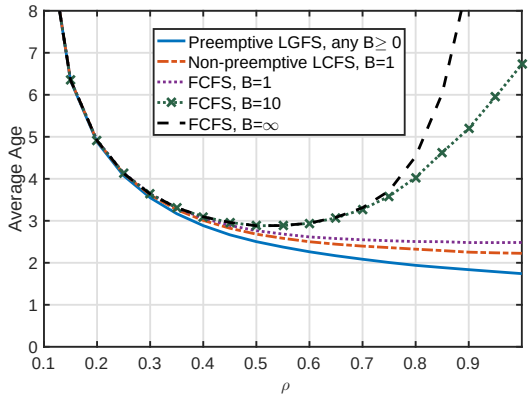

Figure 3: Average age versus traffic intensity ρ for an update system with $m = 1$ server and queue size B.

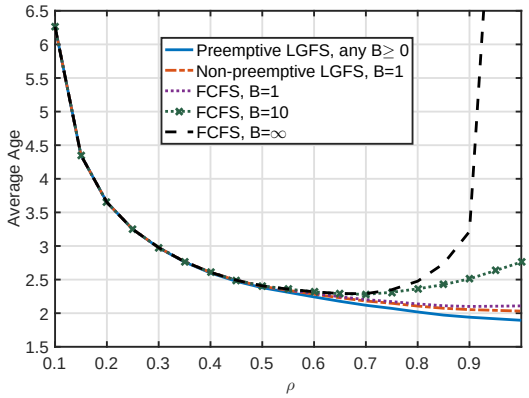

Figure 4: Average age versus traffic intensity ρ for an update system with $m = 5$ servers and queue size B.

generation and arrival times, any given queue size B, $\pi \in \Pi$, and non-decreasing function g,

$$\lim_{T\to\infty} \frac{1}{T}\int_0^T \Delta_{\text{prmp-LGFS}}(t)dt \leq_{\text{st}} \lim_{T\to\infty} \frac{1}{T}\int_0^T \Delta_{\pi}(t)dt,$$

$$\lim_{K\to\infty} \frac{1}{K}\sum_{k=1}^{K} A_{k,\text{prmp-LGFS}} \leq_{\text{st}} \lim_{K\to\infty} \frac{1}{K}\sum_{k=1}^{K} A_{k,\pi},$$

$$\lim_{T\to\infty} \frac{1}{T}\int_0^T g(\Delta_{\text{prmp-LGFS}}(t))dt \leq_{\text{st}} \lim_{T\to\infty} \frac{1}{T}\int_0^T g(\Delta_{\pi}(t))dt,$$

where the limits are assumed to exist.[1]

Hence, the preemptive LGFS policy minimizes time-average age, average peak age, and time-average age penalty for any non-decreasing penalty function g. Finally, the delay and throughput optimality of the preemptive LGFS policy is stated as follows:

Theorem 7. If the packet service times are *i.i.d.* exponentially distributed across time and servers, then for any given packet generation and arrival times, and infinite buffer size $B = \infty$, the preemptive LGFS policy is throughput-optimal and mean-delay-optimal among all policies in Π.

In particular, any work-conserving policy is throughput optimal and mean-delay-optimal. The proof details are provided in Appendix D.

IV. Numerical Results

We present some numerical results to illustrate the age performance of different policies. The packet service times are exponentially distributed with mean $1/\mu = 1$, which is *i.i.d.* across time and servers. The inter-generation times are *i.i.d.* Erlang-2 distribution with mean $1/\lambda$. The number of servers is m. Hence, the traffic intensity is $\rho = \lambda/m\mu$. The queue size is B, which is a non-negative integer.

Figure 3 illustrates the time-average age versus ρ for an information-update system with $m = 1$ server. The time difference between packet generation and arrival $(a_i - s_i)$ is zero, i.e., the update packets arrive in the same order of their generation times. One can observe that the preemptive LGFS policy achieves a better (smaller) age than the FCFS policy analyzed in [4], and the non-preemptive LCFS policy with queue size $B = 1$ [9] which was also named "M/M/1/2*" in [7]. Note that in these prior studies, the time-average age was characterized only for the special case of Poisson arrival process. Moreover, with ordered arrived packets at the server, the LGFS policy and LCFS policy have the same age performance.

Figure 4 plots the time-average age versus ρ for an information-update system with $m = 5$ servers. The time difference between packet generation and arrival, i.e., $a_i - s_i$, is modeled to be either 1 or 100, with equal probability. We found that the age performance of each policy is better than that in Fig. 3, because of the diversity provided by five servers. In addition, the preemptive LGFS policy achieves the best age performance among all plotted policies. It is important to emphasize that the age performance of the preemptive LGFS policy remains the same for any queue size $B \geq 0$. However, the age performance of the non-preemptive LGFS policy and FCFS policy varies with the queue size B when there are multiple servers. We also observe that the average age in case of FCFS policy with $B = \infty$ blows up when the traffic intensity is high. This is due to the increased congestion in the network which leads to a delivery of stale packets. Moreover, in case of FCFS policy with $B = 10$, the average age is high but bounded at high traffic intensity, since the fresh packet has a better opportunity to be delivered in a relatively short period compared with FCFS policy with $B = \infty$. These numerical results validate Theorem 1.

V. Conclusion

In this paper, we considered an information-update system, in which update packets are forwarded to a destination through multiple network servers. It was showed that, if the packet service times are *i.i.d.* exponentially distributed, then for any given arrival process and queue size, the preemptive LGFS pol-

[1]Please refer to Fig. 2 and Section II for the peak age $A_{k,\pi}$.

icy simultaneously optimizes the data freshness, throughput, and delay performance among all causally feasible policies. We will extend these results to more general system settings with general service time distributions.

APPENDIX A
PROOF OF LEMMA 2*

To prove (8), for the sake of contradiction, assume that for some $l \in \{1, \dots, m\}$

$$\hat{\alpha}_{l,P} < \hat{\alpha}_{l,\pi}. \tag{13}$$

Since the arrival process is fixed under all policies, the packet with time stamp $\hat{\alpha}_{l,\pi}$ must have arrived under policy P before the considered time epoch. Let j be the index of the packet with time stamp $\hat{\alpha}_{l,\pi}$. This packet j has four possible cases:

Case 1: The packet j is being processed under both policy P and policy π at the considered time epoch. Without loss of generality, let us assume that $\hat{\alpha}_{n,P} = \hat{\alpha}_{l,\pi}$. We must have $n \leq l$. Otherwise, policy π would have $n - l$ fresh packets that do not exist under policy P, which contradicts the facts that the arrival process is fixed for all policies and that policy P is preemptive LGFS. Hence, we have $\hat{\alpha}_{l,P} \geq \hat{\alpha}_{n,P} = \hat{\alpha}_{l,\pi}$, which contradicts the hypothesis (13).

Case 2: The packet j was preempted by a fresher packet and stored back in the queue under policy P before the considered time epoch, while it is being processed under policy π at the considered time epoch. Thus, the packet j was the oldest packet being processed under policy P before preemption. Hence, we have

$$\hat{\alpha}'_{1,P} = \hat{\alpha}_{l,\pi}, \tag{14}$$

where $\hat{\alpha}'_{1,P}$ denotes the smallest time stamp of the packets being processed by the servers before preemption. After a fresher packet preempted the packet with time stamp $\hat{\alpha}'_{1,P}$, we get

$$\hat{\alpha}_{1,P} \geq \hat{\alpha}_{l,\pi}, \tag{15}$$

where $\hat{\alpha}_{1,P}$ denotes the smallest time stamp of the packets being processed by the servers after preemption. Combining (14) and (15), we obtain

$$\hat{\alpha}_{l,P} \geq \hat{\alpha}_{1,P} \geq \hat{\alpha}_{l,\pi}, \tag{16}$$

which contradicts the hypothesis (13).

Case 3: The packet j or a fresher one is already delivered under policy P, while it is being processed under policy π at the considered time epoch. Hence, we have

$$\hat{U}_P \geq \hat{\alpha}_{l,\pi}. \tag{17}$$

Since $\hat{\alpha}_{l,P} \leq \hat{\alpha}_{l,\pi}$ and $\hat{U}_P \geq \hat{\alpha}_{l,\pi}$, the packet with time stamp $\hat{\alpha}_{l,P}$ is outdated according to policy P. Invoking (4) in the definition of the system state, it must hold that

$$\hat{\alpha}_{l,P} = \hat{U}_P. \tag{18}$$

Combining (17) and (18), we obtain

$$\hat{\alpha}_{l,P} = \hat{U}_P \geq \hat{\alpha}_{l,\pi}, \tag{19}$$

which contradicts the hypothesis (13).

Case 4: The packet j is already delivered under policy π. Hence, we have

$$\hat{U}_\pi = \hat{\alpha}_{l,\pi}. \tag{20}$$

Since $\hat{U}_P \geq \hat{U}_\pi$, the packet j or a fresher one must have been delivered under policy P. Since $\hat{\alpha}_{l,P} \leq \hat{\alpha}_{l,\pi}$ and $\hat{U}_P \geq \hat{U}_\pi = \hat{\alpha}_{l,\pi}$, the packet with time stamp $\hat{\alpha}_{l,P}$ is outdated according to policy P. Invoking (4) in the definition of the system state, it must hold that

$$\hat{\alpha}_{l,P} = \hat{U}_P. \tag{21}$$

Combining (7), (20) and (21), we obtain

$$\hat{\alpha}_{l,P} = \hat{U}_P \geq \hat{U}_\pi = \hat{\alpha}_{l,\pi}, \tag{22}$$

which contradicts the hypothesis (13). As a result of these arguments we have

$$\hat{\alpha}_{i,P} \geq \hat{\alpha}_{i,\pi}, \quad i = 1, \dots, m, \tag{23}$$

which completes the proof.

APPENDIX B
PROOF OF LEMMA 3

Suppose that a packet is delivered in policy P which has a time stamp $\hat{\alpha}_{l,P}$. By coupling, a packet is delivered in policy π at the same time, which has a time stamp $\hat{\alpha}_{l,\pi}$. Hence, we have

$$\hat{U}'_P = \hat{\alpha}_{l,P} \geq \hat{\alpha}_{l,\pi} = \hat{U}'_\pi, \tag{24}$$

which holds for any queue zise $B \geq 0$. From (24) and using Lemma 2*, we obtain

$$\hat{\alpha}'_{i,P} \geq \hat{\alpha}'_{i,\pi}, \quad i = 1, \dots, m. \tag{25}$$

Hence, (10) holds for any queue size $B \geq 0$, which completes the proof.

APPENDIX C
PROOF OF LEMMA 4

Since there is no packet delivery, we have

$$\hat{U}'_P = \hat{U}_P \geq \hat{U}_\pi = \hat{U}'_\pi, \tag{26}$$

From (26) and using Lemma 2*, we obtain

$$\hat{\alpha}'_{i,P} \geq \hat{\alpha}'_{i,\pi}, \quad i = 1, \dots, m. \tag{27}$$

which holds for any queue size $B \geq 0$. Hence, (12) holds for any queue size $B \geq 0$, which completes the proof.

APPENDIX D
PROOF OF THEOREM 7

We follow the same proof technique of Theorem 1. First, we will compare policy P with any work-conserving policy $\pi \in \Pi_{wc}$. For this, we need to define the system state of any policy π:

Definition 8. At any time t, the system state of policy π is specified by $H_\pi(t) = (N_\pi(t), \gamma_\pi(t))$, where $N_\pi(t)$ is the total number of packets in the system at time t. Define $\gamma_\pi(t)$ as the total number of packets that are delivered to the destination at time t. Let $\{H_\pi(t), t \in [0,\infty)\}$ be the state process of policy π, which is assumed to be right-continuous.

To prove Theorem 7, we will need the following lemma.

Lemma 8. For any work-conserving policy $\pi \in \Pi_{wc}$, if $H_P(0^-) = H_\pi(0^-)$, then $\{H_P(t), t \in [0,\infty)\}$ and $\{H_\pi(t), t \in [0,\infty)\}$ are of the same distribution.

Suppose that $\{\hat{H}_P(t), t \in [0,\infty)\}$ and $\{\hat{H}_\pi(t), t \in [0,\infty)\}$ are stochastic processes having the same stochastic laws as $\{H_P(t), t \in [0,\infty)\}$ and $\{H_\pi(t), t \in [0,\infty)\}$, where π is some arbitrary policy in Π_{wc}. Now, we couple the packet delivery times during the evolution of $\hat{H}_P(t)$ to be identical with the packet delivery times during the evolution of $\hat{H}_\pi(t)$. The following two lemmas are needed to prove Lemma 8:

Lemma 9. Suppose that under policy P, $\{\hat{N}'_P, \hat{\gamma}'_P\}$ is obtained by delivering a packet to the destination in the system whose state is $\{\hat{N}_P, \hat{\gamma}_P\}$. Further, suppose that under policy π, $\{\hat{N}'_\pi, \hat{\gamma}'_\pi\}$ is obtained by delivering a packet to the destination in the system whose state is $\{\hat{N}_\pi, \hat{\gamma}_\pi\}$. If

$$\hat{N}_P = \hat{N}_\pi, \hat{\gamma}_P = \hat{\gamma}_\pi,$$

then

$$\hat{N}'_P = \hat{N}'_\pi, \hat{\gamma}'_P = \hat{\gamma}'_\pi. \tag{28}$$

Proof: Since there is a packet delivery, we have

$$\hat{N}'_P = \hat{N}_P - 1 = \hat{N}_\pi - 1 = \hat{N}'_\pi,$$
$$\hat{\gamma}'_P = \hat{\gamma}_P + 1 = \hat{\gamma}_\pi + 1 = \hat{\gamma}'_\pi.$$

Hence, (28) holds, which complete the proof. ∎

Lemma 10. Suppose that under policy P, $\{\hat{N}'_P, \hat{\gamma}'_P\}$ is obtained by adding a new packet to the system whose state is $\{\hat{N}_P, \hat{\gamma}_P\}$. Further, suppose that under policy π, $\{\hat{N}'_\pi, \hat{\gamma}'_\pi\}$ is obtained by adding a new packet to the system whose state is $\{\hat{N}_\pi, \hat{\gamma}_\pi\}$. If

$$\hat{N}_P = \hat{N}_\pi, \hat{\gamma}_P = \hat{\gamma}_\pi,$$

then

$$\hat{N}'_P = \hat{N}'_\pi, \hat{\gamma}'_P = \hat{\gamma}'_\pi. \tag{29}$$

Proof: Because $B = \infty$, no packet is dropped in policy P and policy π. Since there is a new added packet to the system, we have

$$\hat{N}'_P = \hat{N}_P + 1 = \hat{N}_\pi + 1 = \hat{N}'_\pi.$$

Also, there is no packet delivery, hence

$$\hat{\gamma}'_P = \hat{\gamma}_P = \hat{\gamma}_\pi = \hat{\gamma}'_\pi.$$

Thus, (29) holds, which complete the proof. ∎

Proof of Lemma 8: For any sample path, we have that $\hat{N}_P(0^-) = \hat{N}_\pi(0^-)$ and $\hat{\gamma}_P(0^-) = \hat{\gamma}_\pi(0^-)$. This, together with Lemma 9 and 10, implies that

$$\hat{N}_P(t) = \hat{N}_\pi(t), \hat{\gamma}_P(t) = \hat{\gamma}_\pi(t)$$

holds for all $t \in [0,\infty)$. Hence, we have

$$P\left[\hat{H}_P(t) = \hat{H}_\pi(t), t \in [0,\infty)\right] = 1.$$

Because $\{\hat{H}_P(t), t \in [0,\infty)\}$ and $\{H_P(t), t \in [0,\infty)\}$ are of the same distribution, $\{\hat{H}_\pi(t), t \in [0,\infty)\}$ and $\{H_\pi(t), t \in [0,\infty)\}$ are of the same distribution, we get that $\{H_P(t), t \in [0,\infty)\}$ and $\{H_\pi(t), t \in [0,\infty)\}$ are of the same distribution. This completes the proof. ∎

Proof of Theorem 7: As a result of Lemma 8, $\{\gamma_P(t), t \in [0,\infty)\}$ and $\{\gamma_\pi(t), t \in [0,\infty)\}$ are of the same distribution for all $\pi \in \Pi_{wc}$. This implies that all policies $\pi \in \Pi_{wc}$ have the same throughput performance. Also, from Lemma 8, we have that $\{N_P(t), t \in [0,\infty)\}$ and $\{N_\pi(t), t \in [0,\infty)\}$ are of the same distribution for all $\pi \in \Pi_{wc}$. Hence, all policies $\pi \in \Pi_{wc}$ have the same mean-delay performance.

Finally, since the service times are *i.i.d.* across time and servers, service idling only increases the waiting time of the packet in the system. Therefore, the throughput and mean-delay performance under non-work-conserving policies will be worse. As a result, the preemptive LGFS policy is throughput-optimal and mean-delay-optimal among all policies in Π (indeed, any work-conserving policy with infinite queue size $B = \infty$ is throughput optimal and mean-delay-optimal). ∎

REFERENCES

[1] B. Adelberg, H. Garcia-Molina, and B. Kao, "Applying update streams in a soft real-time database system," in *ACM SIGMOD Record*. ACM, 1995, vol. 24, pp. 245–256.

[2] J. Cho and H. Garcia-Molina, "Synchronizing a database to improve freshness," in *Acm Sigmod Record*. ACM, 2000, vol. 29, pp. 117–128.

[3] L. Golab, T. Johnson, and V. Shkapenyuk, "Scheduling updates in a real-time stream warehouse," in *Data Engineering, 2009. ICDE'09. IEEE 25th International Conference on*. IEEE, 2009, pp. 1207–1210.

[4] S. Kaul, R. Yates, and M. Gruteser, "Real-time status: How often should one update?," in *Proc. IEEE INFOCOM*, 2012, pp. 2731–2735.

[5] R. Yates and S. Kaul, "Real-time status updating: Multiple sources," in *Proc. IEEE Int. Symp. Inform. Theory*, July 2012.

[6] L. Huang and E. Modiano, "Optimizing age-of-information in a multi-class queueing system," in *Proc. IEEE Int. Symp. Inform. Theory*, 2015.

[7] M. Costa, M. Codreanu, and A. Ephremides, "Age of information with packet management," in *Proc. IEEE Int. Symp. Inform. Theory*, June 2014, pp. 1583–1587.

[8] N. Pappas, J. Gunnarsson, L. Kratz, M. Kountouris, and V. Angelakis, "Age of information of multiple sources with queue management," in *Proc. IEEE ICC*, June 2015, pp. 5935–5940.

[9] S. Kaul, R. Yates, and M. Gruteser, "Status updates through queues," in *Conf. on Info. Sciences and Systems*, Mar. 2012.

[10] C. Kam, S. Kompella, and A. Ephremides, "Effect of message transmission diversity on status age," in *Proc. IEEE Int. Symp. Inform. Theory*, June 2014, pp. 2411–2415.

[11] M. Costa, S. Valentin, and A. Ephremides, "On the age of channel information for a finite-state markov model," in *Proc. IEEE ICC*, June 2015, pp. 4101–4106.

[12] T. Bacinoglu, E. T. Ceran, and E. Uysal-Biyikoglu, "Age of information under energy replenishment constraints," in *Proc. Info. Theory and Appl. Workshop*, Feb. 2015.
[13] R. Yates, "Lazy is timely: Status updates by an energy harvesting source," in *Proc. IEEE Int. Symp. Inform. Theory*, 2015.
[14] Y. Sun, E. Uysal-Biyikoglu, R. Yates, C. E. Koksal, and N. B. Shroff, "Update or wait: How to keep your data fresh," in *Proc. IEEE INFOCOM*, April 2016.
[15] M. A. Maddah-Ali and D. Tse, "Completely stale transmitter channel state information is still very useful," in *Communication, Control, and Computing (Allerton), 2010 48th Annual Allerton Conference on*, Sept 2010, pp. 1188–1195.
[16] J. Miyaoka and W. Hausman, "How a base stock policy using" stale" forecasts provides supply chain benefits," *Manufacturing & Service Operations Management*, vol. 6, no. 2, pp. 149–162, 2004.
[17] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction*, MIT press, 1998.
[18] M. Shaked and J. G. Shanthikumar, *Stochastic orders*, Springer Science & Business Media, 2007.